Beyond Routing Games: Network Formation Games

Network Games (NG)

- NG model the various ways in which selfish users (i.e., players) strategically interact in using a (either communication, computer, social, etc.) network (modelled as a graph)
- The Internet routing game is a particular type of **network congestion game**
- Other examples of NG: social network games, graphical games, **network design games**, **network creation games**, etc.
- Notice that each of these games is actually a class of games, where each element of the class is specified by the actual input graph, and it is called an instance of the game (i.e, it is a specific game)

Network Formation Games (NFG)

- NFG are NG that aim to capture two competing issues for players when using a network for communication purposes:
  - to minimize the afforded cost
  - to be provided with a high quality of service
- Two big categories of NFG:
  - **Network Design Games** (a.k.a. **Global Connection Games**): Users autonomously design a communication **subnetwork** embedded in an already existing network with the selfish goal of sharing costs in using it for a point-to-point communication
  - **Network Creation Games** (a.k.a. **Local Connection Games**): Users autonomously form *ex-novo* a network that connects them for reciprocal communication (e.g., downloading files in P2P networks, exchanging messages in social networks, etc.)

First case study: Network Design Games (a.k.a. Global Connection Games)

Introduction

- Given a weighted graph G, a **Global Connection Game** (GCG) is a game that models the selfish design of a **communication subnetwork** of G, i.e., a set of point-to-point communication paths, where each path is associated with a player, and the selfish goal of each player is to share the costs for a joint use with other players of the edges on its selected path
- In other words, players:
  - pay for the links they personally use
  - benefit from sharing links with other players in the selected subnetwork

The formal definition of a GCG

- It is given a directed weighted graph $G=(V,E,c)$; c_e will denote the non-negative real weigth of $e \in E$
- k players; each player is associated with a commodity (s_i, t_i) , with $s_i,t_i \in V$, and the strategy for a player i is to select a path P_i in G from s_i to t_i
- Let k_e denote the load of edge e, i.e., the number of players using e; the cost of P_i for player i in a strategy profile $S=(P_1,\ldots,P_k)$ is shared with all the other players using (part of) it, namely:

$$cost_i(S) = \sum_{e \in P_i} c_e/k_e$$

this cost-sharing scheme is called *fair* or *Shapley cost-sharing mechanism*

The formal definition of a GCG (2)

- Given a strategy vector S, the designed network N(S) is given by the union of all paths P_i
- Then, the social-choice function is the utilitarian social cost, namely the total cost of the designed network:

$$C(S)=\sum_i cost_i(S) = \sum_i \sum_{e\in P_i} c_e/k_e = \sum_{e\in N(S)} c_e$$

- Notice that each player has a favorable effect on the cost paid by other players (so-called *cross monotonicity*), as opposed to the *congestion model* of selfish routing

Open questions

- What is a stable network? We use NE as the solution concept, and we will seek for the existence of NE
- How to evaluate the overall quality of a stable network? We compare its cost to that of an optimal (in general, unstable) network, and we will try to estimate a bound on the efficiency loss resulting from selfishness
- Notice that the problem of finding an optimal network is a classic optimization problem (i.e., the network design problem), which is known to be NP-hard even if G is unweighted

Bounding the loss of efficiency

- Remind that a network is *optimal* or *socially efficient* if it minimizes the social cost (i.e., it minimizes the social-choice function)
- We know that the PoA is useful to estimate the loss of efficiency we may have in the worst case, as given by the ratio between the cost of a worst stable network and the cost of an optimal network
- But what about the ratio between the cost of a best stable network and the cost of an optimal network?

The price of stability (PoS)

- **Definition (Schulz & Moses, 2003):** Given a (single-instance) game G and a social-choice function C (which depends on the payoff of all the players), let S be the set of all NE of G. If the payoff represents a cost (resp., a utility) for a player, let OPT be the outcome of G minimizing (resp., maximizing) C. Then, the *Price of Stability (PoS)* of G w.r.t. C is:

$$\mathrm{PoS}_G(C) = \inf_{s \in S} \frac{C(s)}{C(\mathrm{OPT})} \left(\text{resp.}, \sup_{s \in S} \frac{C(s)}{C(\mathrm{OPT})} \right)$$

- **Remark:** If G is a class of games (as for GCG), then its PoS is the maximum/minimum among the PoS of all the instances of G, depending on whether the payoff for a player is either a cost or a utility.

Some remarks

- PoA and PoS are (for positive s.c.f. C)
  - ≥ 1 for minimization (i.e., payoffs are costs) games
  - ≤ 1 for maximization (i.e., payoffs are utilities) games
- PoA and PoS are small when they are close to 1
- PoS is at least as close to 1 than PoA
- In a game with a unique NE, PoA=PoS
- Why studying the PoS?
  - sometimes a nontrivial bound is possible only for PoS
  - PoS quantifies a lower bound to the efficiency loss resulting from selfishness

An example

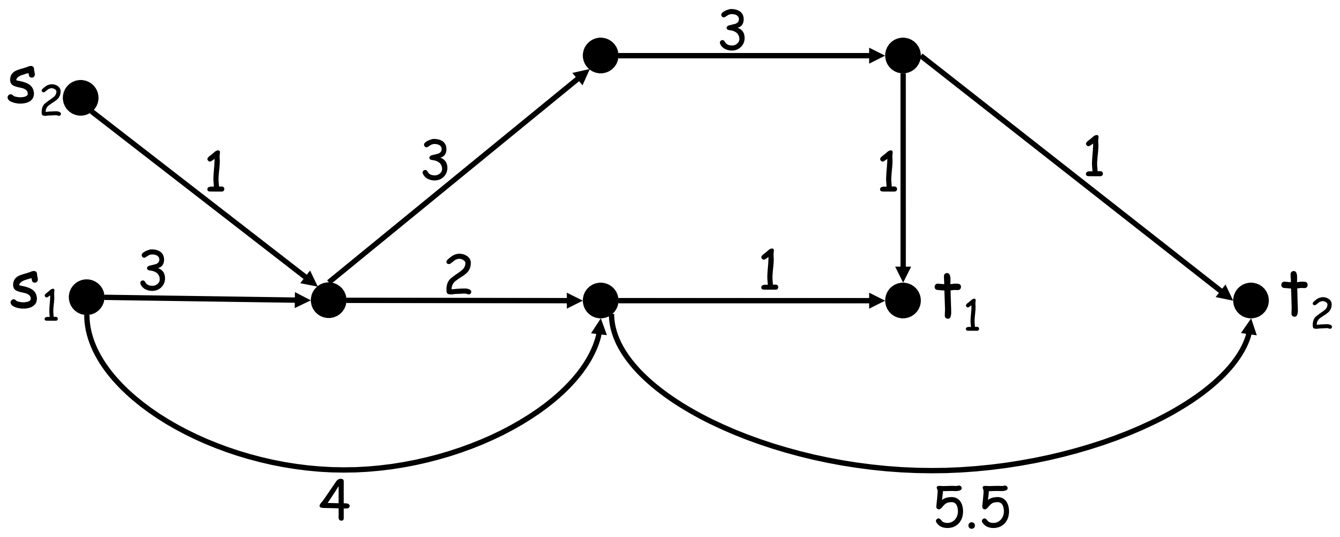

An example

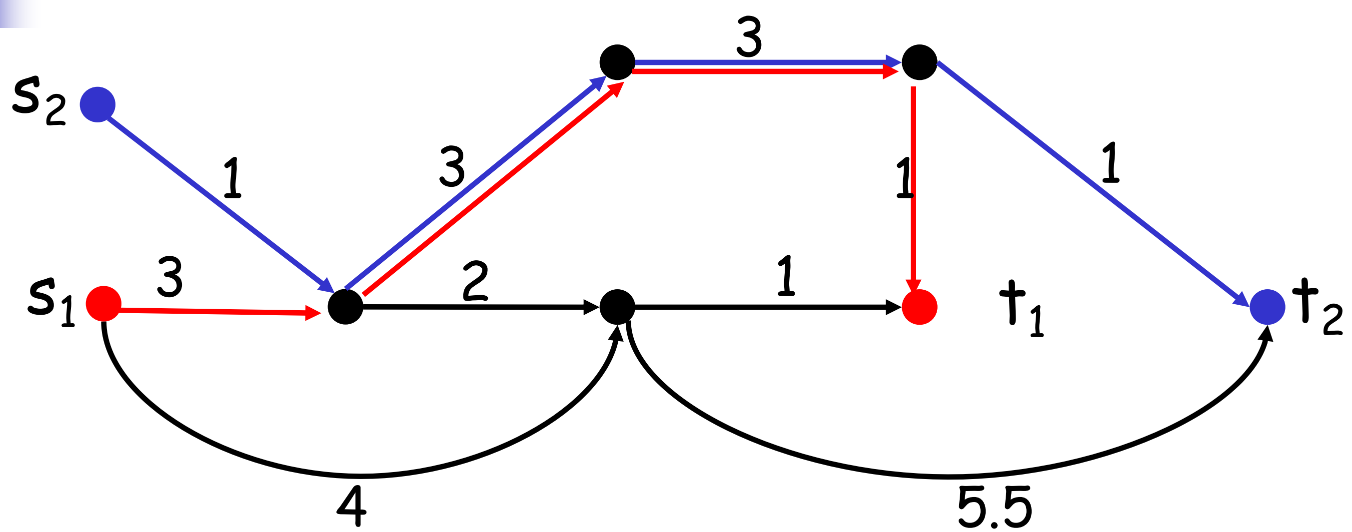

optimal network has cost 12

$cost_1$=7

$cost_2$=5

is it stable?

An example

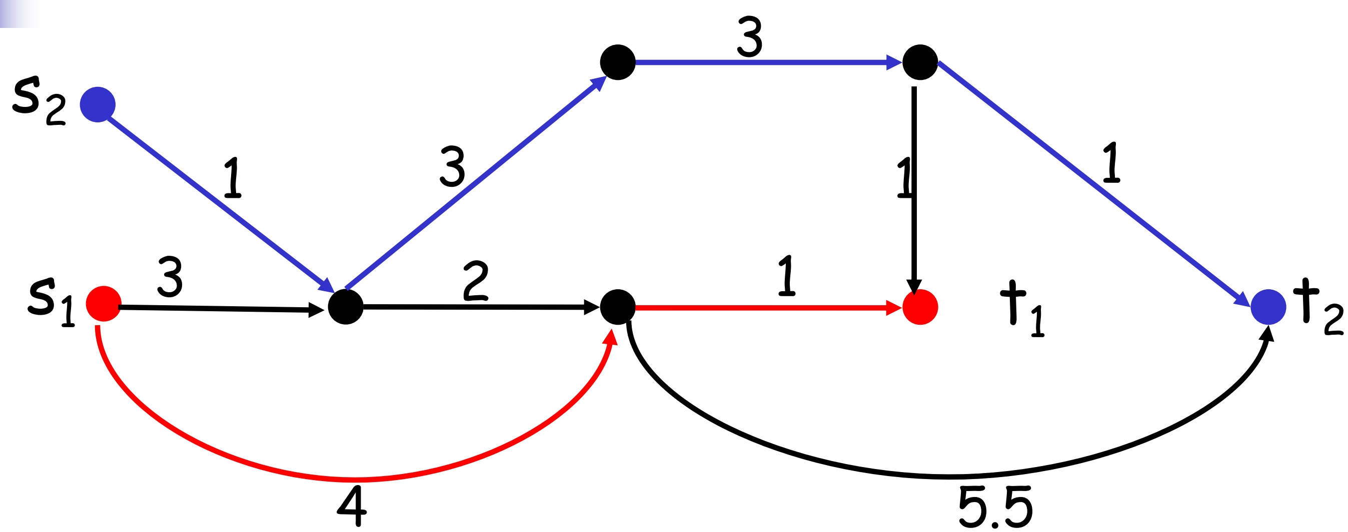

…no!, player 1 can decrease its cost

$cost_1$=5
$cost_2$=8

is it stable? …yes, and has cost 13!

$\Rightarrow$ PoA $\geq$ 13/12, PoS $\leq$ 13/12

An example

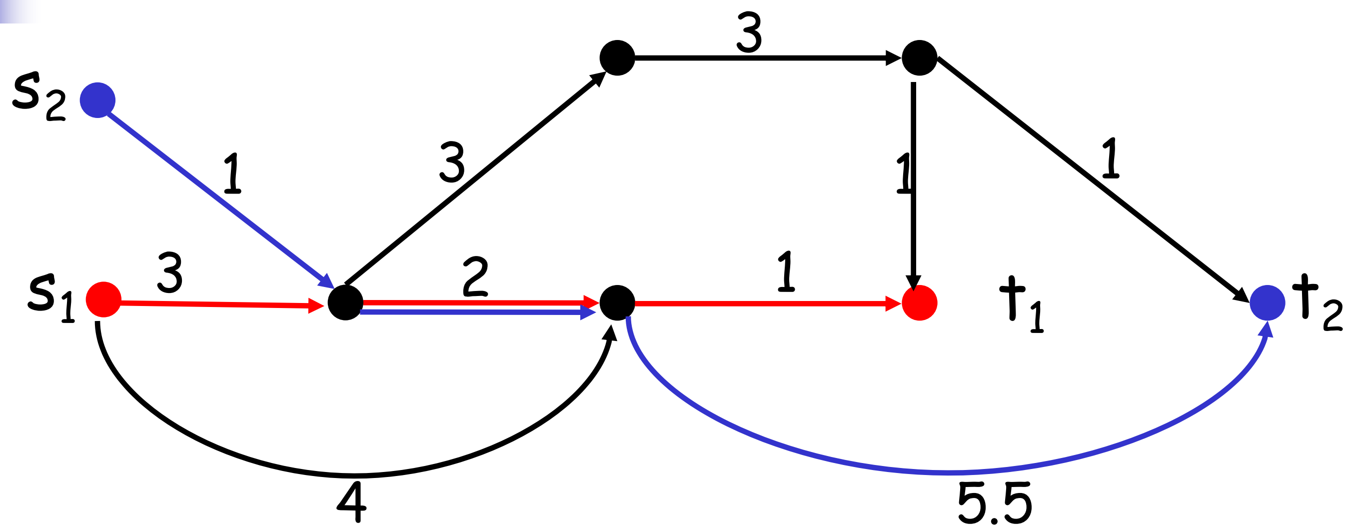

…a best possible NE:

$cost_1$=5
$cost_2$=7.5

the social cost is 12.5 $\Rightarrow$ PoS = 12.5/12

Homework: find a worst possible NE

Theorem 1

Every instance of the GCG has a pure Nash equilibrium, and best response dynamics (i.e., that in which each player at each step selects its best available strategy) always converges.

Theorem 2

The PoA of a GCG with k players is at most k (i.e., **every instance** of the game has PoA $\leq k$), and this is tight (i.e., we can exhibit **an instance** of the game whose PoA is k).

Theorem 3

The PoS of a GCG with k players is at most H_k, the k-th harmonic number (i.e., **every instance** of the game has PoS $\leq H_k$), and this is tight (i.e., we can exhibit **an instance** of the game whose PoS is H_k)

The potential function method

For any finite game, an *exact potential function* Φ is a function that maps every strategy vector S to some real value and satisfies the following condition:

$\forall S=(s_1,\ldots,s_i,,\ldots,s_k)$, let $s'_i \neq s_i$, and let $S'=(s_1,\ldots,s'_i,\ldots,s_k)$, then

$$\Phi(S)-\Phi(S') = cost_i(S)-cost_i(S').$$

A (finite) game that does possess an exact potential function is called *potential game*

Lemma 1

Every potential game has at least one pure Nash equilibrium, namely the strategy vector $\hat{S}$ that **minimizes** (resp., **maximizes**) Φ, assuming players' payoffs are costs (resp., utilities).

Proof (minimization): Observe that Φ is bounded. Then, starting from $\hat{S}=(\hat{s}_1,\ldots,\hat{s}_i,\ldots,\hat{s}_k)$, consider any move by a player i that results in a new strategy vector $S=(\hat{S}_{-i},s_i) = (\hat{s}_1,\ldots,\hat{s}_{i-1},s_i,\ldots,\hat{s}_k)$. Since $\Phi(\hat{S})$ is minimum, we have:

$$\underbrace{\Phi(\hat{S})-\Phi(S)}_{\leq 0} = cost_i(\hat{S})-cost_i(S)$$

$\Rightarrow$ $cost_i(\hat{S}) \leq cost_i(S)$ $\Rightarrow$ player i cannot decrease its cost, thus $\hat{S}$ is a NE. ∎

Convergence in potential games

Observation: any state S with the property that $\Phi(S)$ cannot be decreased by changing any single component of S is a NE by the same argument. Furthermore, by definition, improving moves for players decrease the value of the potential function, which is **bounded**. Together, these observations imply the following result.

Lemma 2

In any finite potential game, best response dynamics always converges to a Nash equilibrium

☹ However, it may be the case that converging to a NE takes an exponential (in the number of players) number of steps!

…turning our attention to the global connection game…

Let Ψ be the following function mapping any strategy vector S to a real value [Rosenthal 1973]:

$$\Psi(S) = \Sigma_{e \in N(S)} \Psi_e(S)$$

where (recall that k_e is the number of players using e)

$$\Psi_e(S) = c_e \cdot H_{k_e} = c_e \cdot (1+1/2+\ldots+1/k_e).$$

Lemma 3 (Ψ is a potential function)

Let $S=(P_1,\dots,P_k)$, let P'_i be an alternative path for some player i, and define a new strategy vector $S'=(S_{-i},P'_i)$. Then:

$$\Psi(S) - \Psi(S') = cost_i(S) - cost_i(S').$$

Proof:

When player i switches from P_i to P'_i, some edges of N(S) increase their load by 1, some others decrease it by 1, and the remaining do not change it. Then, it suffices to notice that:

- If load of edge e increases by 1, its contribution to the potential function increases by $c_e/(k_e+1)$
- If load of edge e decreases by 1, its contribution to the potential function decreases by c_e/k_e

$$\Rightarrow \Psi(S) - \Psi(S') = \Psi(S) - \Psi(S-P_i+P'_i) = \Psi(S) - \left(\Psi(S) - \sum_{e\in P_i} c_e/k_e + \sum_{e\in P'_i} c_e/(k_e+1)\right) = cost_i(S) - cost_i(S').$$

Existence of a NE

Theorem 1

Every instance of the GCG has a pure Nash equilibrium, and best response dynamics always converges.

Proof: From Lemma 3, a GCG is a potential game, and from Lemma 1 and 2 best response dynamics converges to a pure NE. ■

☺ It can be shown that finding a best response for a player is polynomial (it suffices to find a shortest path in G where each edges e is weighted as $c_e/(k_e+1)$)

☹ Instead, it can be shown that finding a NE of cost at most C (and so, finding a best/worst NE) is NP-hard!

Price of Anarchy: a lower bound

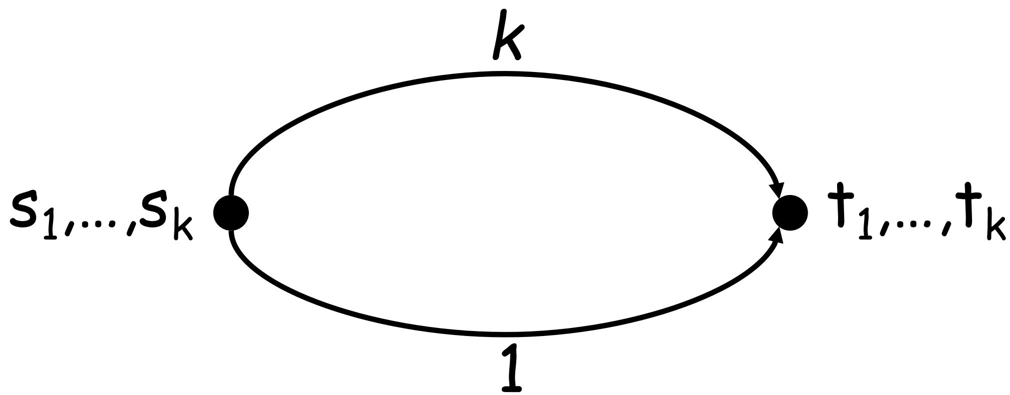

optimal network has cost 1

best NE: all players use the lower edge ⇒ PoS is 1 ☺

worst NE: all players use the upper edge ⇒ PoA is k ☹

Upper-bounding the PoA

Theorem 2

The price of anarchy in the global connection game with k players is at most k (and so, from the previous lower bound, this is tight).

Proof: Let $OPT=(P_1^*,\dots,P_k^*)$ denote the optimal set of paths (i.e., a set of paths minimizing C). Let Π_i be a shortest path in $G=(V,E,c)$ between s_i and t_i w.r.t. c, and let $c(\Pi_i) = \Sigma_{e\in\Pi_i} c_e$ be the length of such a path. Finally, let S be any NE. Observe that $cost_i(S)\leq c(\Pi_i)$ (otherwise the player i would change to Π_i). Then:

$$C(S) = \sum_{i=1}^{k} cost_i(S) \leq \sum_{i=1}^{k} c(\Pi_i) \leq \sum_{i=1}^{k} c(P_i^*) =$$

$$\sum_{i=1}^{k} \sum_{e\in P_i^*} c_e \leq \sum_{i=1}^{k} \sum_{e\in P_i^*} k\cdot c_e/k_e = \sum_{i=1}^{k} k\cdot cost_i(OPT) = k\cdot C(OPT).$$

PoS for GCG: a lower bound

ε>0: small value

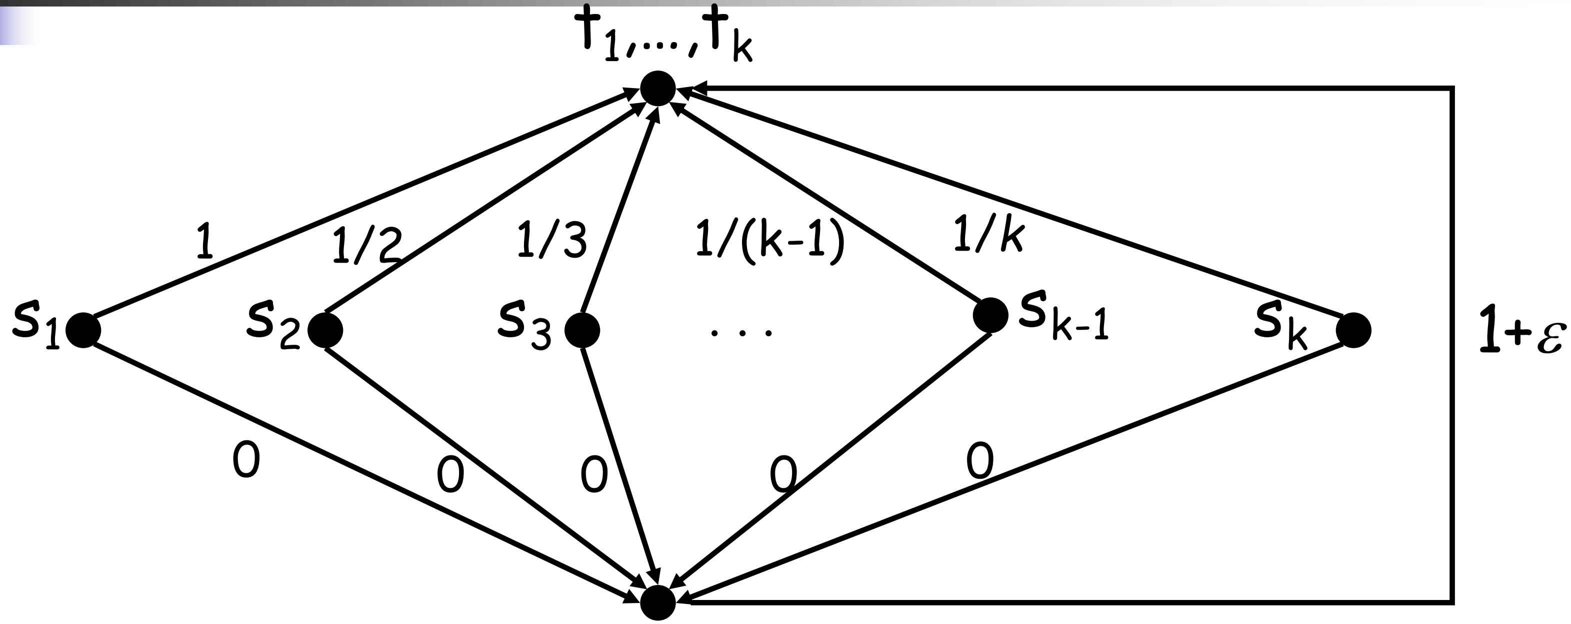

PoS for GCG: a lower bound

$\varepsilon > 0$: small value

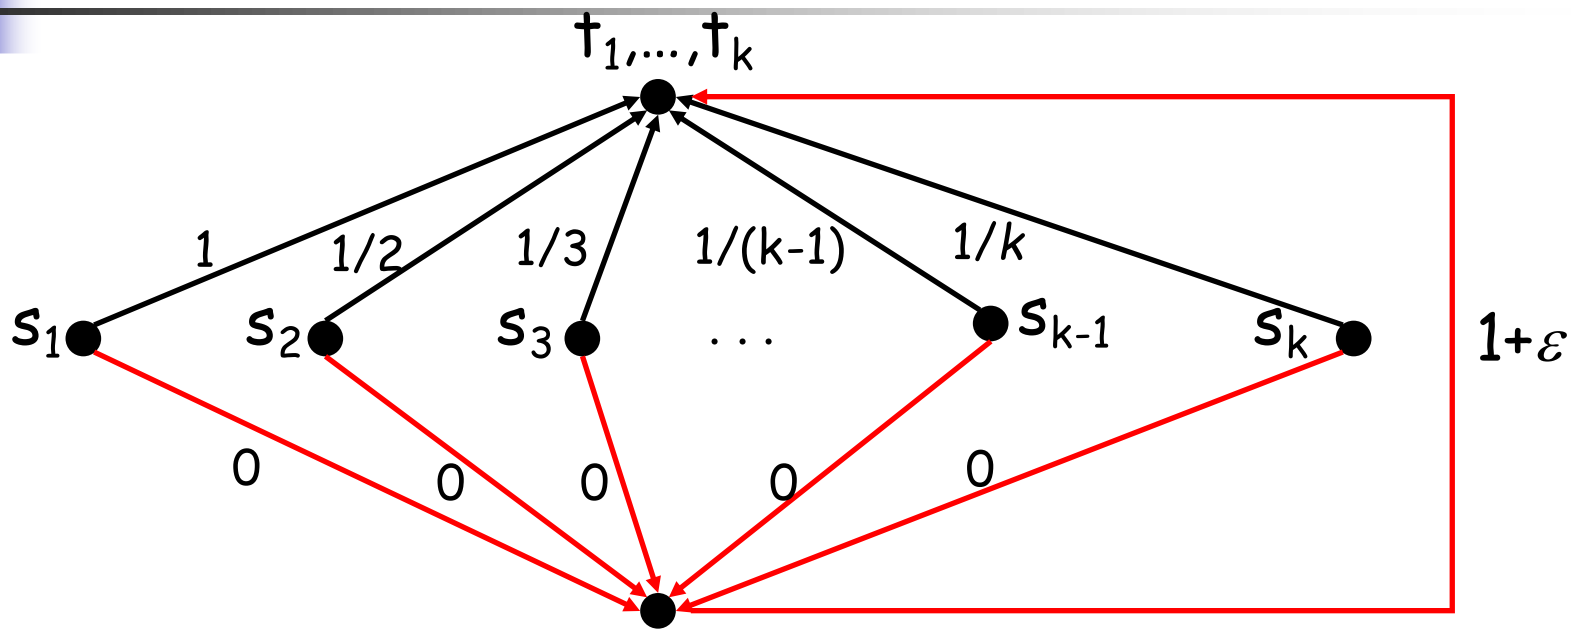

The optimal solution has a cost of $1+\varepsilon$

is it stable?

PoS for GCG: a lower bound

ε>0: small value

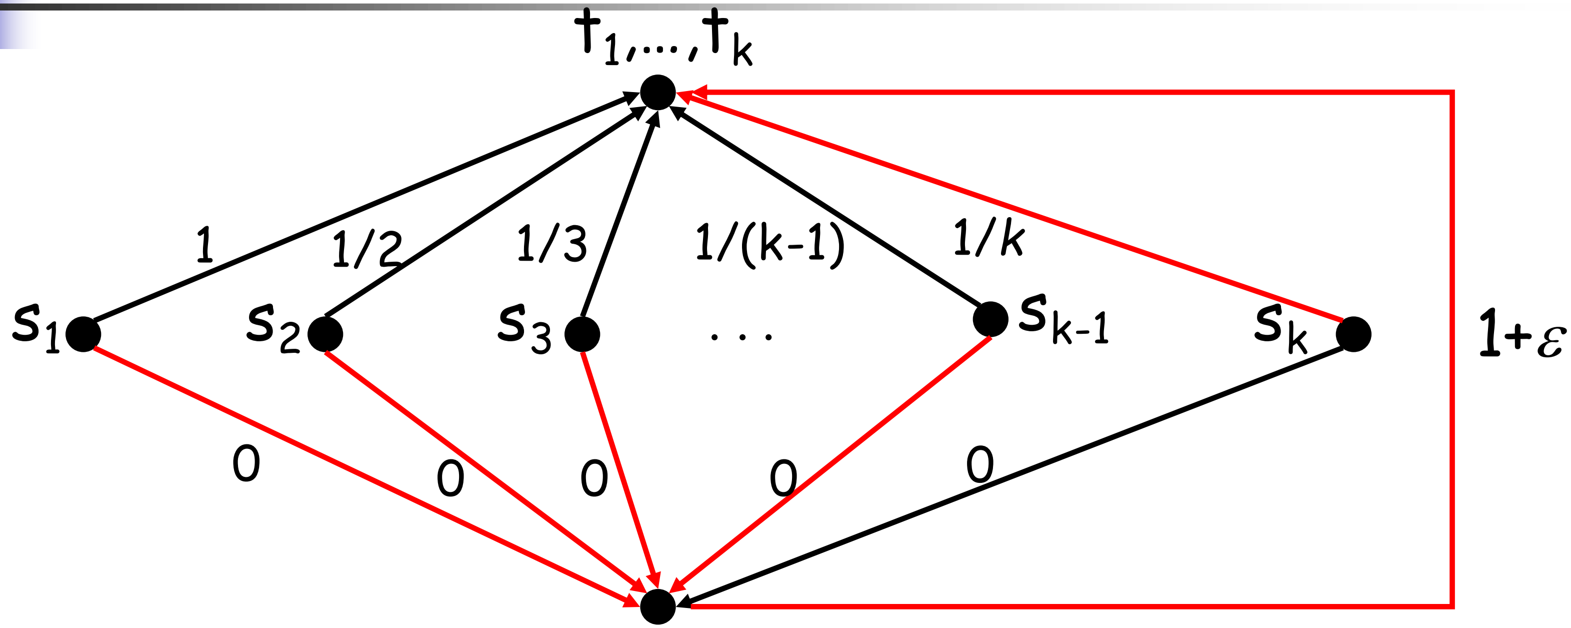

…no! player k can decrease its cost…

is it stable?

PoS for GCG: a lower bound

$\varepsilon>0$: small value

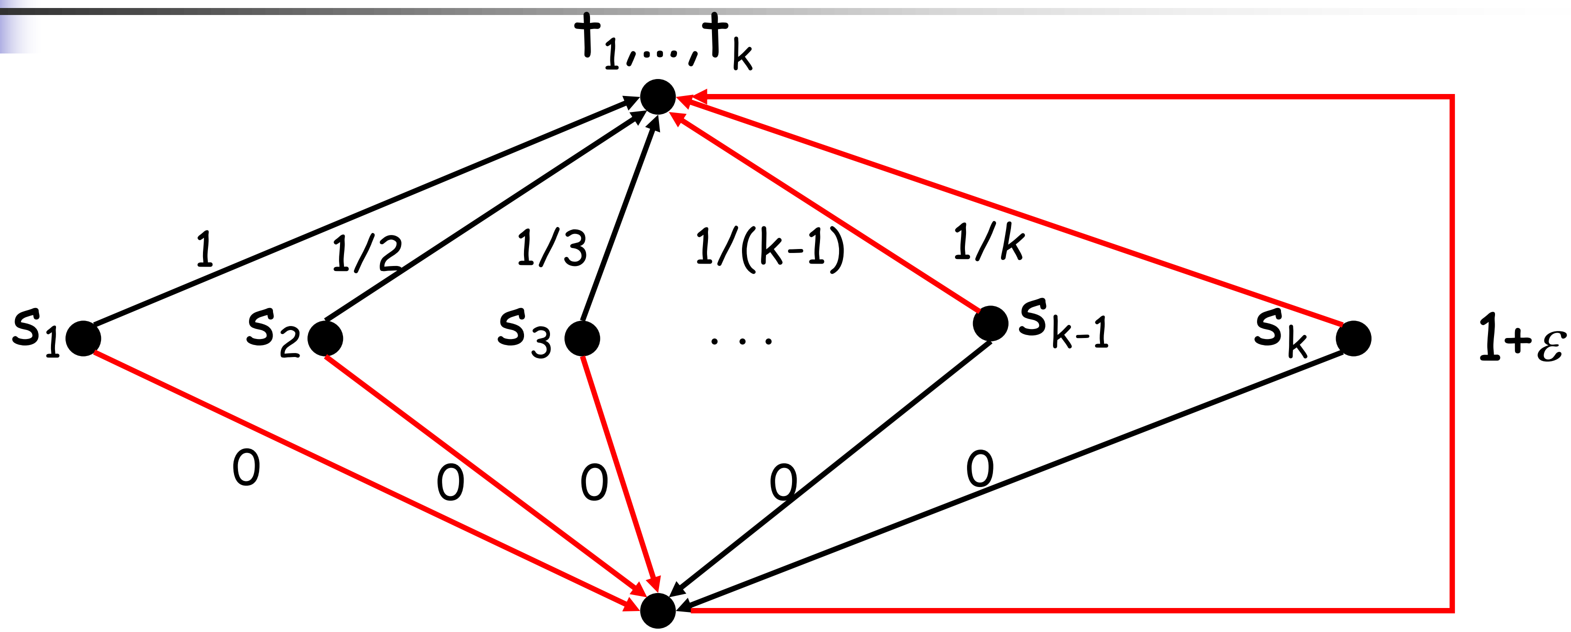

…no! player k-1 can decrease its cost…

is it stable?

PoS for GCG: a lower bound

$\varepsilon > 0$: small value

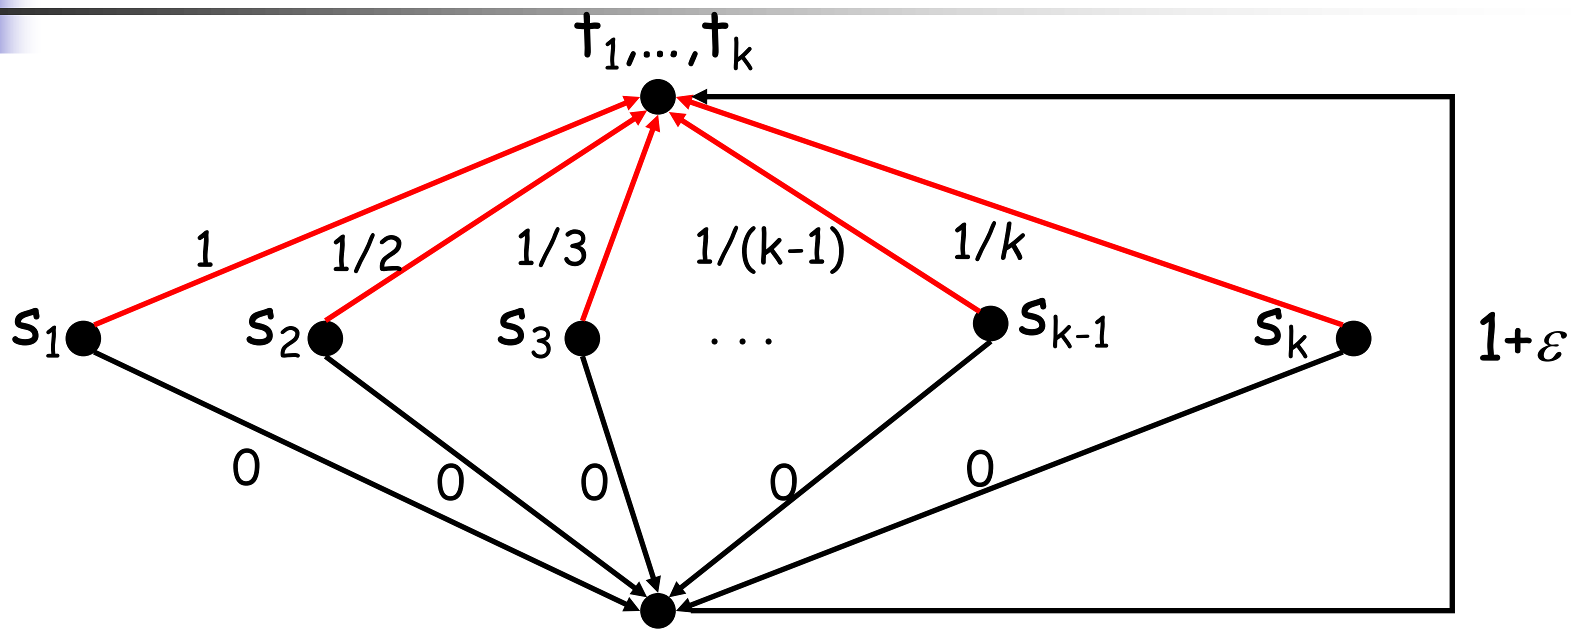

The only stable network

social cost: $C(S) = \sum_{j=1}^{k} 1/j = H_k \leq \ln k + 1$ k-th *harmonic number*

Lemma 4

Suppose that we have a potential game with potential function Φ, and assume that for any outcome S we have

$$C(S)/A \le \Phi(S) \le B\, C(S)$$

for some A,B>0. Then the price of stability is at most AB.

Proof:

Let $\hat{S}$ be the strategy vector minimizing Φ (i.e., $\hat{S}$ is a NE, from Lemma 1). Let S* be the strategy vector minimizing the social cost

we have:

$$C(\hat{S})/A \le \Phi(\hat{S}) \le \Phi(S^*) \le B\, C(S^*)$$

$$\Rightarrow PoS \le C(\hat{S})/C(S^*) \le A \cdot B.$$

Lemma 5 (Bounding Ψ)

For any strategy vector S in the GCG, we have:

$$C(S) \leq \Psi(S) \leq H_k C(S).$$

Proof: Indeed:

$$\Psi(S) = \Sigma_{e \in N(S)} \Psi_e(S) = \Sigma_{e \in N(S)} c_e \cdot H_{k_e}$$

$$\Rightarrow \Psi(S) \geq C(S) = \Sigma_{e \in N(S)} c_e$$

and $\Psi(S) \leq H_k \cdot C(S) = \Sigma_{e \in N(S)} c_e \cdot H_k$.

■

Upper-bounding the PoS

Theorem 3

The price of stability in the global connection game with k players is at most H_k, the k-th harmonic number (and so, from the previous lower bound, this is tight).

Proof: From Lemma 3, a *GCG* is a potential game, and from Lemma 5 and Lemma 4 (with *A*=1 and B=H_k), its PoS is at most H_k.

■